arXiv:1702.02264v1 [stat.AP] 8 Feb 2017

Precision Therapeutic Biomarker Identification with Application to the Cancer Genome Project

Hongmei Liu and J. Sunil Rao

Division of Biostatistics, University of Miami, Miami, FL 33136
h.liu7@med.miami.edu
JRao@biostat.med.miami.edu

Cancer cell lines have frequently been used to link drug sensitivity and resistance with genomic profiles. To capture genomic complexity in cancer, the Cancer Genome Project (CGP) (Garnett et al., 2012) screened 639 human tumor cell lines with 130 drugs ranging from known chemotherapeutic agents to experimental compounds. Questions of interest include: i) can cancer-specific therapeutic biomarkers be detected, ii) can drug resistance patterns be identified along with predictive strategies to circumvent resistance using alternate drugs, iii) can biomarkers of drug synergies be predicted ? To tackle these questions, following statistical challenges still exist: i)biomarkers cluster among the cell lines; ii) clusters can overlap (e.g. a cell line may belong to multiple clusters); iii) drugs should be modeled jointly. We introduce a multivariate regression model with a latent overlapping cluster indicator variable to address above issues. A generalized finite mixture of multivariate regression (FMMR) model in connection with the new model and a new EM algorithm for fitting are proposed. Re-analysis of the dataset sheds new light on the therapeutic inter-relationships between cancers as well existing and novel drug behaviors for the treatment and management of cancer.

Key Words: Cancer biomarkers; EM algorithm; finite mixture of multivariate regression model; LASSO; overlapping clustering.

1 Introduction

1.1 Data Description

The use of drugs to selectively target specific genetic alterations in defined patient subpopulations has seen significant successes. One example can be found in the treatment of chronic myeloid leukaemia (CML) where the first consistent chromosomal abnormality associated with a human cancer was identified back in the 1960s. Fast forward to the 1980s where the consequence of this abnormality was discovered to be the production of an abnormal gene called BCR-ABL. Intense drug discovery programs were initiated to shut down the activity of BCR-ABL, and in 1992, imatinib (Gleevec) was developed. In 1998, the drug was tested in CML patients who had exhausted standard treatment options and whose life expectancy was limited, with remarkable results in their blood counts returning to normal. In 2001, the FDA approved imatinib. Today, a once commonly fatal cancer now has a five-year survival rate of 95% (Druker et al., 2006).

Achievements like this largely inspire today's high throughput screening studies of linking cancer drugs (known or in development) to specific genomic changes which could be used as therapeutic biomarkers. The hope is that such analyses will shed light on biological mechanisms underlying drug sensitivity, tumor resistance and potential drug combination synergies.

Cancer cell lines have frequently been used as a convenient way of conducting such studies. For a systematic search of therapeutic biomarkers to a variety of cancer drugs, the Cancer Genome Project (CGP) (Garnett et al., 2012) screened 639 human tumor cell lines, which represent much of the tissue-type and genetic diversity of human cancers, with 130 drugs. These drugs, including approved drugs, drugs in development as well as experimental tool compounds, cover a wide range of targets and processes involved in cancer biology. A range of 275–507 cell lines were screened for each drug. The effect of a 72h drug treatment on cell viability was examined to derive such measures of drug sensitivity as the half-maximal inhibitory concentration (IC_{50}). The cell lines underwent sequencing of 64 known cancer genes, genome-wide analysis of copy number gains and losses, and expression profiling of 14,500 genes.

Given the degree of complexity of this dataset, the multivariate analysis of variance (MANOVA) and the Elastic-Net regression applied in Garnett et al. (2012) are insufficient for precise knowledge discovery. First, the marginal drug-feature associations discovered in MANOVA rarely reflect true relationships, as it is more likely that sensitivity of cancer cells to drugs depends on a multiplicity of genomic and epigenomic features with potential interactions. Second, the Elastic-Net regression fails to concern following issues: 1) since the 639 cell lines come from a variety of cancer tissue types, there is likely additional heterogeneity manifested as subpopulations with overlaps in data; 2) note that the 130 drugs (response variables) are hardly independent, one can improve the prediction accuracy by modeling with multiple drugs (Breiman and Friedman, 1997).

Moreover, there are some direct questions of interest from a subject matter perspective that we want to address. These include, i) can cancer-specific therapeutic biomarkers be detected, ii) can drug resistance patterns be identified along with predictive strategies to circumvent resistance using alternate drugs, iii) can biomarkers of combination therapies be identified to help predict synergies in drug activities ? To tackle these questions and previously discussed statistical challenges, we propose a multivariate regression model with a latent overlapping cluster indicator variable. Fitting procedures inducing concurrent variable selection are introduced.

The rest of the paper is organized as follows. In Section 1.2, we give a selective overview of existing clustering and overlapping clustering methods for general (without response variables) and regression data. In Section 2, a new statistical model is introduced, a generalized FMMR model in connection with the new model and a new EM algorithm for fitting are provided. We also establish a type of consistency optimality for estimation of the generalized FMMR model and perform some small simulation studies to empirically demonstrate this. In Section 3, we put forward another fitting solution to the new model for comparison with the generalized FMMR model. Section 4 contains a comprehensive re-analyses of the CGP data using the proposed method. Discussion is included in Section 5.

1.2 Relevant Statistical Literature Review

Clustering is a well established technique to group data elements based on a measure of similarity. Traditional clustering techniques generate partitions so that each data point belongs to one and only one cluster. It has long been recognized that such ideal partition seldom exists in real data (Needham, 1965). It is more likely that clusters overlap in some

parts. To handle the overlapping issue, Lazzeroni and Owen (2002) put forward the well-known plaid model for two-sided overlapping clustering for gene expression data, see also Turner et al. (2005) for improved plaid model and Zhang (2010) for Bayesian plaid model formulation. Its numerical solution, however, produces unsatisfactory cluster retrievals (see Section 3). Other overlapping clustering methods include the "naive" finite mixture (FM) model with a hard threshold on posterior membership probabilities, the probabilistic model (Banerjee et al., 2005) and the multiplicative mixture model based approach (Fu and Banerjee, 2008).

More often interest centers on investigating the relationship between response variables $\mathbf{Y} \in \mathbb{R}^q$ and covariates $\mathbf{X} \in \mathbb{R}^p$ by fitting a regression model rather than to explore the $\mathbf{X}$ or $\mathbf{Y}$ on its own. In regression analyses with $q = 1$, finite mixture of regression (FMR) models are commonly used to capture unobserved cross-sectional heterogeneity in the data (Jedidi et al., 1996). The FMR model postulates that a sample of observations come from a finite mixture of latent partitioning sub-populations with each sub-population represented by a regression model. It was first introduced to statistical literatures by Quandt (1972). DeSarbo and Cron (1988) proposed an EM algorithm (Dempster et al., 1977) based maximum likelihood estimation for the FMR model. Khalili and Chen (2012) put forward a penalized likelihood approach for variable selection in FMR models. As a natural generalization to multivariate responses case, Jones and McLachlan (1992) introduced the FMMR model. Grun and Leisch (2008) proposed a R package *flexmix* for fitting FMMR models, however it assumes that the response variables are independent.

2 The Proposed Method

2.1 Statistical modeling

For a sample (of cell lines) of n observations, denote $\mathbf{y}_i = (y_{i1}, \ldots, y_{iq})^T \in \mathbb{R}^q$ a vector of responses (IC_{50} values), $\mathbf{x}_i = (x_{i1}, \ldots, x_{ip_n})^T \in \mathbb{R}^{p_n}$ a vector of predictors (genomic markers) and $\varepsilon_i = (\varepsilon_{i1}, \ldots, \varepsilon_{iq})^T$ a vector of random errors for the ith observation. Assume $\varepsilon_i \overset{iid}{\sim} N_q(\mathbf{0}, \mathbf{\Sigma})$ for $i = 1, \ldots, n$. The following proposed model allows overlapping clustering for multivariate regression data,

$$\mathbf{y}_i = \sum_{k=1}^{K} \mathbf{B}_k^T \mathbf{x}_i P_{ik} + \varepsilon_i, \quad i = 1, \ldots, n, \tag{1}$$

where K is the total number of clusters, $\mathbf{B}_k$ is an unknown $p_n \times q$ coefficient matrix for the kth cluster, and $P_{ik} \in \{0, 1\}$ is 1 if observation i belongs to the kth cluster, otherwise 0. In traditional clustering problem, it assumes that each observation belongs to exactly one cluster, namely $\sum_{k=1}^{K} P_{ik} = 1$ for all i. Therefore we allow $\sum_{k=1}^{K} P_{ik} \geq 1$ so that each observation can belong to multiple clusters.

We provide some interpretation for the clusters in model (1). A cluster k contains a subset of observations for whom $P_{ik} = 1$. We postulate that not all genomic features are relevant in describing the cluster k, thus assume a *sparse* coefficient matrix $\mathbf{B}_k$. Since the sparse patterns can vary with k, each cluster is represented by a unique set of biomarkers. For observations belong to multiple clusters, their response variables are explained by multiple sets of biomarkers, indicative of involving in several biological processes simultaneously.

2.2 The FMMR model

When $\sum_{k=1}^{K} P_{ik} = 1$, model (1) can be characterized by a hierarchical structure which ends up with an FMMR model. Consider a latent cluster membership random variable z_i for observation i from (1). Given $\sum_{k=1}^{K} P_{ik} = 1$, the range of z_i equals to $\{1, \ldots, K\}$. Assume that $P(z_i = k) = \pi_k$ for each k, then $\sum_{k=1}^{K} \pi_k = 1$ and $\pi_k \geq 0$. Response $\mathbf{y}_i$ from model (1) satisfies

$$(\mathbf{y}_i \mid z_i = k, \mathbf{x}_i, \mathbf{B}_k, \mathbf{\Sigma}) \sim N_q(\mathbf{B}_k^T \mathbf{x}_i, \mathbf{\Sigma}). \tag{2}$$

Denote $\mathbf{\Theta} = (\mathbf{B}_1, \ldots, \mathbf{B}_K, \mathbf{\Sigma}, \pi)$ with $\pi = (\pi_1, \ldots, \pi_{K-1})^T$. We can derive the joint density of $\mathbf{y}_i$ and z_i as

$$f(\mathbf{y}_i, z_i \mid \mathbf{x}_i, \mathbf{\Theta}) = \prod_{k=1}^{K} \{\pi_k f(\mathbf{y}_i \mid z_i = k, \mathbf{x}_i, \mathbf{B}_k, \mathbf{\Sigma})\}^{I(z_i = k)}. \tag{3}$$

Summarizing (3) over $z_i = 1, \ldots, K$ yields the FMMR model

$$f(\mathbf{y}_i \mid \mathbf{x}_i, \mathbf{\Theta}) = \sum_{k=1}^{K} \pi_k f(\mathbf{y}_i \mid z_i = k, \mathbf{x}_i, \mathbf{B}_k, \mathbf{\Sigma}). \tag{4}$$

When using the EM algorithm to estimate model (4), $(\mathbf{x}_i, \mathbf{y}_i)$ is regarded as incomplete data for missing z_i, so one works on the complete joint density in (3). For a sample of n observations from (1), the complete log-likelihood function of $\mathbf{\Theta}$ becomes

$$l_n(\mathbf{\Theta}) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_{ik} \log \pi_k + \sum_{i=1}^{n} \sum_{k=1}^{K} z_{ik} \log f(\mathbf{y}_i \mid z_i = k, \mathbf{x}_i, \mathbf{B}_k, \mathbf{\Sigma}), \tag{5}$$

where $z_{ik} = I(z_i = k)$.

Since we do not expect all genomic markers to be informative, variable selection, to obtain a parsimonious model is necessary. Khalili and Chen (2012) and Khalili and Lin (2013) (for diverging model size) introduced a penalized likelihood approach for variable selection in FMR models, which was shown consistent in variable selection and highly efficient in computation. We define a penalized complete log-likelihood function as

$$\tilde{l}_n(\mathbf{\Theta}) = l_n(\mathbf{\Theta}) - \rho_n(\mathbf{\Theta}), \tag{6}$$

where

$$\rho_n(\mathbf{\Theta}) = \sum_{k=1}^{K} \pi_k \rho_{nk}(\mathbf{B}_k).$$

The following two penalty functions are currently used to meet different demands in variable selection:

1) the L_1-penalty in LASSO (Tibshirani, 1996) for simultaneous estimation and variable selection

$$\rho_{nk}(\mathbf{B}_k) = \lambda_k \|\mathbf{B}_k\|_1$$

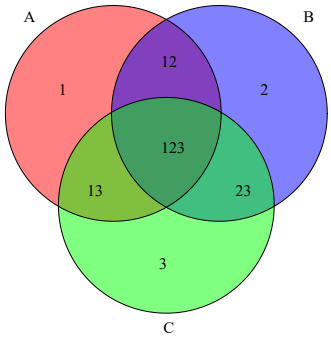

Figure 1: Overall overlapping patterns given $K = 3$.

2) a linear combination of the L_1- and L_2-penalty in Elastic-Net (Zou and Hastie, 2005) for simultaneous estimation and selection of grouped features

$$\rho_{nk}(\mathbf{B}_k) = \lambda_{k1}\|\mathbf{B}_k\|_1 + \lambda_{k2}\|\mathbf{B}_k\|_2^2,$$

where $\|\cdot\|_1$ and $\|\cdot\|_2$ are respectively the L_1- and L_2-norms and the tuning parameters $\lambda_k, \lambda_{k1}, \lambda_{k2} \geq 0$. The L_1-penalty is singular at the origin, thus can shrink some coefficients to exact 0 for sufficiently large λ_k or λ_{k1} (Fan and Li, 2001).

Their revised EM algorithm can easily be generalized to maximize $\tilde{l}_n(\boldsymbol{\Theta})$ in (6), hence we bypass here. In next section, generalized FMMR model and EM algorithm are devised to fit overlapping multivariate regression data.

2.3 The Generalized FMMR Model

Finite mixture models (including FMR and FMMR) can only fit partitions, as can be seen from Section 2.2. To enable overlapping clustering, one often uses a "naive" approach by applying a hard threshold α to finite mixture models. This approach assigns a sample to cluster k if its posterior probability of belonging to cluster k is larger than a pre-specified α, hence enabling a sample to belong to multiple clusters. As pointed out by Banerjee et al. (2005), this method is problematic because it is not a natural generative model for overlapping clustering, since one underlying assumption of the finite mixture model is that each observation comes from one and only one mixture component. In this section however, we introduce a generalized FMMR model to retrieve overlapping clusters when $\sum_{k=1}^{K} P_{ik} \geq 1$. This generalized model retains its ability in fitting partitions.

Suppose we have K *objective clusters* indexed by 1 to K. These objective clusters refer to clusters defined in (1) and can overlap with each other, resulting in $2^K - 1$ types of overlapping patterns. Let K equal to 3 for an example, overall overlapping patterns are then composed of $S = \{1, 2, 3, 12, 13, 23, 123\}$. Figure 1 shows the particular part each overlapping pattern in S represents. By definition, overlapping patterns in S are mutually exclusive. And each observation i from (1) belongs to one and only one overlapping pattern.

In the next, we define $2^K - 1$ *hypothetical clusters*. Each hypothetical cluster (simply called "cluster" since after) represents an overlapping pattern defined in advance and is

indexed by an element in T,

$$T = \cup_{s=1}^{K} \{(l_1 \dots l_s) : \{l_1, \dots, l_s\} \subseteq \{1, \dots, K\}\}.$$

Cluster $(l_1 \dots l_s)$ implies its members belonging to objective clusters $l_1, \dots, l_s$.

We now introduce a latent cluster membership random variable z_i for observation i from model (1) and characterize the model via a hierarchical structure. Given $\sum_{k=1}^{K} P_{ik} \geq 1$, the range of z_i becomes T. Further define

$$P(z_i = (l_1 \dots l_s)) = \pi_{(l_1 \dots l_s)}. \tag{7}$$

Then vector $\pi = (\pi_{(l_1 \dots l_s) \in T})^T$ has

$$\sum_{(l_1 \dots l_s) \in T} \pi_{(l_1 \dots l_s)} = 1, \quad \pi_{(l_1 \dots l_s)} \geq 0.$$

Response $\mathbf{y}_i$ from (1) satisfies

$$(\mathbf{y}_i \mid z_i = (l_1 \dots l_s), \mathbf{x}_i, \{\mathbf{B}_{l_k}\}_{k=1}^{s}, \mathbf{\Sigma}) \sim N_q \left(\sum_{k=1}^{s} \mathbf{B}_{l_k}^T \mathbf{x}_i, \mathbf{\Sigma} \right). \tag{8}$$

Let $\mathbf{\Theta} = (\mathbf{B}_1, \dots, \mathbf{B}_K, \mathbf{\Sigma}, \pi)$. By (7) and (8), the joint density of $\mathbf{y}_i$ and z_i equals

$$f(\mathbf{y}_i, z_i \mid \mathbf{x}_i, \mathbf{\Theta}) = \prod_{(l_1 \dots l_s) \in T} \left\{ \pi_{(l_1 \dots l_s)} f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l_k}\}_{k=1}^{s}, \mathbf{\Sigma}\right) \right\}^{I(z_i = (l_1 \dots l_s))}. \tag{9}$$

Summarizing (9) over z_i leads to the generalized FMMR model

$$f(\mathbf{y}_i \mid \mathbf{x}_i, \mathbf{\Theta}) = \sum_{(l_1 \dots l_s) \in T} \pi_{(l_1 \dots l_s)} f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l_k}\}_{k=1}^{s}, \mathbf{\Sigma}\right). \tag{10}$$

Note that if $\pi_{(l_1 \dots l_s)} = 0$ for $s > 1$, (10) reduces to the traditional FMMR model in (4).

Then the (conditional) log-likelihood function of $\mathbf{\Theta}$ for a sample of n observations from (1) is

$$l_n^0(\mathbf{\Theta}) = \sum_{i=1}^{n} \log \left(\sum_{(l_1 \dots l_s) \in T} \pi_{(l_1 \dots l_s)} f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l_k}\}_{k=1}^{s}, \mathbf{\Sigma}\right) \right).$$

Maximizing above ordinary likelihood function yields non-zero estimates for all regression coefficients. To induce variable selection and remove noise predictors from the regression model, we propose a penalized log-likelihood function

$$\tilde{l}_n^0(\mathbf{\Theta}) = l_n^0(\mathbf{\Theta}) - \sum_{k=1}^{K} \sum_{\substack{(l_1 \dots l_s): \\ k \in \{l_1, \dots, l_s\}}} \pi_{(l_1 \dots l_s)} \rho_{nk}(\mathbf{B}_k), \tag{11}$$

where $\rho_{nk}(\mathbf{B}_k)$ is the LASSO or Elastic-Net penalty function in Section 2.2. Notation $\sum_{\substack{(l_1\ldots l_s):\\ k\in\{l_1,\ldots,l_s\}}}$ indicates that it summarizes over $(l_1\ldots l_s) \in T$ for which $k \in \{l_1,\ldots,l_s\}$. The penalty imposed on $\mathbf{B}_k$ is proportional to $\sum_{\substack{(l_1\ldots l_s):\\ k\in\{l_1,\ldots,l_s\}}} \pi_{(l_1\ldots l_s)}$, which by definition is proportional to the number of observations involved in the kth objective cluster. This is a strategy, similar to Khalili and Chen (2012), of relating the penalty to sample sizes for enhanced power of the method.

2.4 Numerical Solution to the Generalized FMMR Model

We use the renowned EM algorithm for optimization of the generalized FMMR model. The complete log-likelihood function of $\boldsymbol{\Theta}$ is

$$\begin{aligned} l_n(\boldsymbol{\Theta}) =& \sum_{i=1}^{n} \sum_{(l_1\ldots l_s)\in T} z_{i,(l_1\ldots l_s)} \log \pi_{(l_1\ldots l_s)} \\ &+ \sum_{i=1}^{n} \sum_{(l_1\ldots l_s)\in T} z_{i,(l_1\ldots l_s)} \log f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l_k}\}_{k=1}^{s}, \boldsymbol{\Sigma}\right), \end{aligned}$$

where $z_{i,(l_1\ldots l_s)} = I(z_i = (l_1\ldots l_s))$.

The penalized complete log-likelihood function for concurrent variable selection is then

$$\tilde{l}_n(\boldsymbol{\Theta}) = l_n(\boldsymbol{\Theta}) - \sum_{k=1}^{K} \sum_{\substack{(l_1\ldots l_s):\\ k\in\{l_1,\ldots,l_s\}}} \pi_{(l_1\ldots l_s)} \rho_{nk}(\mathbf{B}_k). \tag{12}$$

Due to the overlap setting, each $\mathbf{B}_k$ involves in multiple clusters. Therefore coefficient matrix $\mathbf{B}_k$ for $k = 1,\ldots,K$ can not be optimized independently like in usual EM algorithm. Instead we sequentially update $\mathbf{B}_k$ given the rest are known. Specifically take $\rho_{nk}(\mathbf{B}_k)$ as a LASSO penalty function for an example, our revised EM algorithm iteratively maximizes $\tilde{l}_n(\boldsymbol{\Theta})$ in two steps:

E-step: Given $\boldsymbol{\Theta}^m$, estimate $z^m_{i,(l_1\ldots l_s)}$ with its posterior probability by applying the Bayes' rule to (9) and (10),

$$\hat{P}(z^m_{i,(l_1\ldots l_s)} = 1 \mid \mathbf{y}_i, \mathbf{x}_i, \boldsymbol{\Theta}^m) = \frac{\pi^m_{(l_1\ldots l_s)} f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l_k}\}_{k=1}^{s}, \boldsymbol{\Sigma}^m\right)}{\sum_{(l'_1\ldots l'_s)\in T} \pi^m_{(l'_1\ldots l'_s)} f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l'_k}\}_{k=1}^{s}, \boldsymbol{\Sigma}^m\right)}.$$

M-step: Given $\mathbf{Z}^m = (z^m_{i,(l_1\ldots l_s)})$, update the mixing proportions π^{m+1} by

$$\pi^{m+1}_{(l_1\ldots l_s)} = \frac{1}{n}\sum_{i=1}^{n} z^m_{i,(l_1\ldots l_s)}, \quad (l_1\ldots l_s) \in T.$$

Note that this is obtained by maximizing the leading term of (12) with respect to π. This simplified updating scheme however works well in simulation studies.

Given $\mathbf{Z}^m$, π^{m+1} and $\mathbf{\Sigma}^m$, sequentially update $\mathbf{B}_k^{m+1} \mid \{\mathbf{B}_s^m\}_{s=1,s\neq k}^{s=K}$ by

$$\mathbf{B}_k^{m+1} = \arg\max_{\mathbf{B}_k} \sum_{i=1}^{n} \sum_{\substack{(l_1\ldots l_s):\\ k\in\{l_1,\ldots,l_s\}}} z_{i,(l_1\ldots l_s)}^m \log f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l_j}^m\}_{j=1,l_j\neq k}^{s}, \mathbf{\Sigma}^m\right) - \sum_{\substack{(l_1\ldots l_s):\\ k\in\{l_1,\ldots,l_s\}}} \pi_{(l_1\ldots l_s)}^{m+1}\lambda_k\|\mathbf{B}_k\|_1. \quad (13)$$

By (8),

$$\mathbf{B}_k^{m+1} = \arg\min_{\mathbf{B}_k} \sum_{\substack{(l_1\ldots l_s):\\ k\in\{l_1,\ldots,l_s\}}} tr\big((\mathbf{Y}^* - \mathbf{X}\mathbf{B}_k)^T\mathbf{z}_{(l_1\ldots l_s)}^m(\mathbf{Y}^* - \mathbf{X}\mathbf{B}_k)(\mathbf{\Sigma}^m)^{-1}\big) + 2\sum_{\substack{(l_1\ldots l_s):\\ k\in\{l_1,\ldots,l_s\}}} \pi_{(l_1\ldots l_s)}^{m+1}\lambda_k\|\mathbf{B}_k\|_1, \quad (14)$$

where

$$\mathbf{y}_i^* = \mathbf{y}_i - \sum_{l\in\{l_1,\ldots,l_s\}\setminus k} (\mathbf{B}_l^m)^T\mathbf{x}_i, \quad \mathbf{Y}^* = (\mathbf{y}_1^*, \ldots, \mathbf{y}_q^*),$$

$$\mathbf{z}_{(l_1\ldots l_s)}^m = diag\left(z_{1,(l_1\ldots l_s)}^m, \ldots, z_{n,(l_1\ldots l_s)}^m\right).$$

Above (14) is a multivariate regression problem with a LASSO penalty for estimation, which can be solved by the MRCE algorithm (Rothman et al., 2010). If we ignore the covariance structure of $\mathbf{\Sigma}^m$ merely in (14), estimation of $\mathbf{B}_k^{m+1}$ reduces to q independent LASSO regression problems. In this case, more complex penalty functions such as the Elastic-Net and fused LASSO penalties can easily be applied. Therefore we also have investigated the performance of this simplified strategy in simulation studies. It works surprisingly comparable to original method utilizing the MRCE algorithm for estimation of $\mathbf{B}_k^{m+1}$.

Given $\mathbf{Z}^m$ and $\{\mathbf{B}_k^{m+1}\}_{k=1}^K$, update $\mathbf{\Sigma}^{m+1}$ by

$$\begin{aligned}\mathbf{\Sigma}^{m+1} &= \arg\max_{\mathbf{\Sigma}} \sum_{i=1}^{n} \sum_{(l_1\ldots l_s)\in T} z_{i,(l_1\ldots l_s)}^m \log f\left(\mathbf{y}_i \mid z_i, \mathbf{x}_i, \{\mathbf{B}_{l_k}^{m+1}\}_{k=1}^{s}, \mathbf{\Sigma}\right)\\ &= \arg\min_{\mathbf{\Sigma}} \sum_{i=1}^{n} \sum_{(l_1\ldots l_s)\in T} -z_{i,(l_1\ldots l_s)}^m \log|\mathbf{\Sigma}^{-1}|\\ &\quad + \sum_{(l_1\ldots l_s)\in T} tr\big((\mathbf{Y} - \mathbf{X}\sum_{k=1}^{s}\mathbf{B}_{l_k}^{m+1})^T\mathbf{z}_{(l_1\ldots l_s)}^m(\mathbf{Y} - \mathbf{X}\sum_{k=1}^{s}\mathbf{B}_{l_k}^{m+1})\mathbf{\Sigma}^{-1}\big). \quad (15)\end{aligned}$$

By taking the derivative of (15) according to $\mathbf{\Sigma}^{-1}$, we get

$$\mathbf{\Sigma}^{m+1} = \frac{\sum_{(l_1\ldots l_s)\in T}(\mathbf{Y} - \mathbf{X}\sum_{k=1}^{s}\mathbf{B}_{l_k}^{m+1})^T\mathbf{z}_{(l_1\ldots l_s)}^m(\mathbf{Y} - \mathbf{X}\sum_{k=1}^{s}\mathbf{B}_{l_k}^{m+1})}{\sum_{(l_1\ldots l_s)\in T}\sum_{i=1}^{n} z_{i,(l_1\ldots l_s)}^m},$$

where $\sum_{(l_1 \ldots l_s) \in T} \sum_{i=1}^{n} z^m_{i,(l_1 \ldots l_s)} = n$.

Although we can also penalize the inverse covariance matrix $\mathbf{\Sigma}^{-1}$ of (15) like in Rothman et al. (2010), simulations (not presented in the paper) show that penalizing $\mathbf{\Sigma}^{-1}$ results in a lower degree of clustering accuracy than not penalizing $\mathbf{\Sigma}^{-1}$. This may be because by penalizing $\mathbf{\Sigma}^{-1}$ one introduces bias into its estimation, a biased estimate of $\mathbf{\Sigma}^{-1}$ in return deteriorates the estimate of $z_{i,(l_1 \ldots l_s)}$ in E-step. Thus we choose to not penalize $\mathbf{\Sigma}^{-1}$.

Commencing with an initial value $\mathbf{\Theta}^0$, the algorithm iterates between E- and M-steps until the relative change in log-likelihood, $|\big(l_n^{m+1}(\hat{\mathbf{\Theta}}) - l_n^m(\hat{\mathbf{\Theta}})\big)/l_n^m(\hat{\mathbf{\Theta}})|$, is smaller than some threshold value, taken as 10^{-5} in simulation studies and 10^{-3} in real data analysis. Additionally, a cluster, whose mixing proportion is smaller than some threshold value taken as 0.01 in the paper, will be removed during iterations to avoid over estimations.

2.5 Selection of Tuning Parameters and the K

In preceding penalized likelihood approach, one needs to choose the values of component-wise tuning parameters λ_k for $k = 1, \ldots, K$, which controls the complexity of an estimated model. The data-driven method cross-validation (CV) (Stone, 1977) is frequently adopted in literatures. Here we use a component-wise 10-fold CV method for tuning parameters selection in (14).

Moreover, selection of the number of components K is essential in finite mixture models (including FMR and FMMR). In applications, the choice can be based on prior knowledge of data analysts. With respect to formatted methodologies, information criteria (IC) remains by far the most popular strategy for selection of K. See Claeskens et al. (2008) for general treatments on this topic. Here we choose the K minimizing below IC_n,

$$IC_n(K) = -2l_n(\mathbf{\Theta}) + N_K a_n. \tag{16}$$

Above N_k is the effective number of parameters in the model,

$$N_K = |\{\mathbf{B}_k, k = 1, \ldots, K\}| + |\pi| - 1 + |\mathbf{\Sigma}|,$$

where $|A|$ calculates the number of *nonzero* elements in A. In (16), a_n is a positive sequence depending on n. The well known AIC (Akaike, 1974) and BIC (Schwarz et al., 1978) correspond to $a_n = 2$ and $a_n = \log(n)$ respectively. It was shown in Keribin (2000) that under general regularity conditions, BIC can identify the true order of a finite mixture model asymptotically. Feasibility of his results to FMR or FMMR models is unknown yet. We examined the performance of BIC for selection of K in a generalized FMMR model via simulation studies. It obtained a high degree of accuracy.

2.6 The Asymptotic Properties and Simulation Studies

Khalili and Lin (2013) showed that their approach by maximizing a penalized log-likelihood function for estimation of the FMR model is consistent in both estimation and variable selection under certain regularity conditions. Denote $f(\mathbf{w}; \mathbf{\Theta})$ the joint density function of data $\mathbf{w} = (\mathbf{x}, \mathbf{y})$ with $\mathbf{\Theta} \in \mathbf{\Omega}$. The conditional density function of $\mathbf{y}$ given $\mathbf{x}$ follows an FMR model in Khalili and Lin (2013) and a generalized FMMR model in (10) in our paper. As defined in (11), the penalized log-likelihhood function is a summation of the $\log(f(\mathbf{w}_i; \mathbf{\Theta}))$'s

minus a penalty function. Because the regularity conditions for asymptotic establishments in Khalili and Lin (2013) are made on $f(\mathbf{w};\boldsymbol{\Theta})$ and the penalty function directly, their theoretical achievements can be extended to our problem.

Denote $\mathbf{B}_k^j$ the jth column of $\mathbf{B}_k$ for $j = 1,\ldots,q$. Let $\mathbf{B}_k^j = (\mathbf{B}_{k1}^j, \mathbf{B}_{k2}^j)$ to divide the coefficient vector into non-zero and zero subsets. Denote $\boldsymbol{\Theta}^0$ the true value of $\boldsymbol{\Theta}$, $\boldsymbol{\Theta}^0 = (\boldsymbol{\Theta}_1^0, \boldsymbol{\Theta}_2^0)$ is the corresponding decomposition such that $\boldsymbol{\Theta}_2^0$ contains all zero coefficients, and $\hat{\boldsymbol{\Theta}}_n = (\hat{\boldsymbol{\Theta}}_{n1}, \hat{\boldsymbol{\Theta}}_{n2})$ is its estimate. Since the dimension of $\hat{\boldsymbol{\Theta}}_{n1}$ increases with n, we investigate the asymptotic distribution of its finite linear transformation, $\mathbf{D}_n\hat{\boldsymbol{\Theta}}_{n1}$, where $\mathbf{D}_n$ is an $l \times d_{n1}$ constant matrix with a finite l and d_{n1} is the dimension of $\boldsymbol{\Theta}_1^0$. Moreover $\mathbf{D}_n\mathbf{D}_n^T \to \mathbf{D}$ and $\mathbf{D}$ is a positive definite and symmetric matrix. We assume that K is independent of the sample size n and known beforehand. Its selection is discussed in Section 2.5.

Lemma 1. *Suppose the penalty function $\rho_{nk}(\mathbf{B}_k)$ satisfies conditions $\mathcal{P}_0 - \mathcal{P}_2$ and the joint density function $f(\mathbf{w};\boldsymbol{\Theta})$ satisfies conditions $R_1 - R_5$ in Khalili and Lin (2013).*

1. If $\frac{p_n^4}{n} \to 0$, then there exists a local maximizer $\hat{\boldsymbol{\Theta}}_n$ for the penalized log-likelihood $\tilde{l}_n(\boldsymbol{\Theta})$ with

$$\|\hat{\boldsymbol{\Theta}}_n - \boldsymbol{\Theta}^0\| = O_p\left\{\sqrt{\frac{p_n}{n}}(1 + q_{2n})\right\},$$

where $q_{2n} = \max_{kij}\left\{\frac{|\rho'_{nk}(b_{kij}^0)|}{\sqrt{n}} : b_{kij}^0 \neq 0\right\}$, b_{kij}^0 is an entry of $\mathbf{B}_k^0$.

2. If $\rho_{nk}(\mathbf{B}_k)$ also satisfies condition $\mathcal{P}_3$ in Khalili and Lin (2013) and $\frac{p_n^5}{n} \to 0$, for any $\sqrt{n/p_n}$-consistent maximum likelihood estimate $\hat{\boldsymbol{\Theta}}_n$, as $n \to \infty$ it has:
i. Variable selection consistency:

$$P(\hat{\mathbf{B}}_{k2}^j = \mathbf{0}) \to 1, \quad k = 1,\ldots,K \text{ and } j = 1,\ldots,q.$$

ii. Asymptotic normality:

$$\sqrt{n}\mathbf{D}_n\mathcal{I}_1^{-1/2}(\boldsymbol{\Theta}_1^0)\left\{\left[\mathcal{I}_1(\boldsymbol{\Theta}_1^0) - \frac{\rho_n''(\boldsymbol{\Theta}_1^0)}{n}\right](\hat{\boldsymbol{\Theta}}_{n1} - \boldsymbol{\Theta}_1^0) + \frac{\rho_n'(\boldsymbol{\Theta}_1^0)}{n}\right\} \to^d N(0, \mathbf{D}), \tag{17}$$

where $\mathcal{I}_1(\boldsymbol{\Theta}_1^0)$ is the Fisher information matrix under the true subset model.

Proof. Write $\boldsymbol{\Theta} = (\theta_1, \theta_2, \ldots, \theta_{t_n})^T$, where t_n is the total number of parameters in the model. Then Lemma 1 is a direct extension from Khalili and Lin (2013). □

By Lemma 1, $\hat{\boldsymbol{\Theta}}_n$ under the LASSO or Elastic Net penalty has a convergence rate $\sqrt{p_n/n}$ via appropriate choice of the tuning parameters. However consistent estimation does not necessarily guarantee consistent variable selection. By the Lemma, the LASSO or Elastic Net penalty does not lead to consistent variable selection, because on one hand λ_k must be large enough to achieve sparsity, on the other hand the bias term q_{2n} is proportional to the $\lambda_k/\sqrt{n}$. This problem can be solved by using adaptive LASSO or adaptive Elastic Net penalty instead, which leads to concurrent estimation and variable selection consistency.

We conducted a sequence of small simulation studies to numerically demonstrate optimality properties of the new method. They are presented in Web Appendix A.

3 Generalization of the Plaid Model

Lazzeroni and Owen (2002) proposed the plaid model to decompose a gene expression data as the sum of overlapping layers (clusters). Each layer contains a subset of genes and samples, while each gene and sample can participate in multiple layers. Denote x_{ij} a gene expression data entry and K is the total number of layers. The plaid model is

$$x_{ij} = \omega_0 + \sum_{k=1}^{K} (\omega_k + \alpha_{ik} + \beta_{jk}) P_{ik} Q_{jk},$$

where $\omega_0, \omega_k, \alpha_{ik}, \beta_{jk} \in \mathbb{R}$, P_{ik} is 1 if gene i is in the kth gene-bolck and otherwise 0, Q_{jk} is 1 if sample j is in the kth sample-block and otherwise 0. An iterative algorithm was put forward to consecutively search for the layers.

In fact, model (1) is a generalized plaid model by introducing covariates into the model so that it can cluster overlapping multivariate regression data. Therefore, estimation procedures of the plaid model can also be used to estimate (1). In parallel to the plaid model, the parameters are estimated by minimizing the Q,

$$Q = \frac{1}{2} \sum_{i=1}^{n} \|\mathbf{y}_i - \sum_{k=1}^{K} \mathbf{x}_i^T \mathbf{B}_k P_{ik}\|_2^2 + \sum_{k=1}^{K} \rho_{nk}(\mathbf{B}_k),$$

where $\rho_{nk}(\mathbf{B}_k)$ is the penalty function defined in Section 2.2.

Detailed procedures for optimizing Q are shown in Algorithm 1 in Web Appendix B, which is a generalization of the improved iterative algorithm by Turner et al. (2005) to multivariate regression. At each time, the algorithm searches for a layer that explains as much of the rest data as possible; once a layer has been found, it is subtracted from the data and remains unchanged; the algorithm stops when no further layers can be found. However this algorithm introduces bias to parameter estimations due to its discrete updating scheme, where previous recruited layers can not be refined as new layers are recruited into the model. To address this issue, we further revised Algorithm 1 to enable joint optimization of all layers. Detailed procedures are presented in Algorithm 2 in Web Appendix B.

4 Therapeutic Biomarker Identification for the CGP Data

We now turn our attention back to analyses of the CGP high throughput drug sensitivity dataset for cancer described in Section 1.1. An updated version of this data (http://www.cancerrxgene.org/downloads/) contains 707 human tumor cell lines as samples, 140 drugs as response variables, and 13831 genomic features as covariates (including the tissue type, rearrangements, mutation status of 71 cancer genes, continuous copy number data of 426 genes causally implicated in cancer, as well as genome-wide transcriptional profiles). The response variables, made up with IC_{50} values from pairwise drug-cell-line screening, have some missing values. Therefore the data was first filtered by removing cell lines for which less than 50% of the drugs were tested, resulting in 591 cell lines remaining. In the remaining data, about 37.4% of the cell lines are with missing values. These missing values were then imputed via the a random forest imputation algorithm (Ishwaran et al.,

2008).

Direct Q1: Can cancer-specific therapeutic biomarkers be detected?

We applied our new method to identify patterns of cancer-specific therapeutic biomarkers. Note that by "cancer-specific" we do not assume separate patterns for each cancer type but rather clusters that may be *driven* by one or a small number of cancer types. Throughout the analysis, we fixed K to a value of 3, although as previously shown, a BIC model selection approach could also be used. However, fixing the value of K does not limit the interesting findings that we can still find and saves on computational time.

Due to the scale and complexity of the data, the analysis was conducted in three steps. Although simulation studies show that modeling with multiple response variables yields much higher clustering accuracy than modeling with a single response variable, it is unreasonable to simply fit all 140 drugs (responses) in one generalized FMMR model. Because by doing so, one assumes that the cell line assignments to clusters are the same for all 140 drugs, which can hardly be true. Thus in our analysis, we first divided the 140 drugs into several groups and then fitted each drug group with a generalized FMMR model, in which to capture active grouped genomic features, the Elastic-Net penalty along with a simplified updating scheme of $\mathbf{B}_k$ discussed in Section 2.4 was used.

In the first step, each drug c was fitted by a generalized FMMR model, from which we got drug-specific cell line assignments, $\hat{\mathbf{Z}}_c$, and cluster-wise coefficient estimates, $\hat{\mathbf{B}}_k^c$ for $k = 1, \ldots, 3$. In the second step, we used the affinity-propagation clustering (APC) algorithm (Frey and Dueck, 2007) to group the 140 drugs based on results from step 1. The grouping was conducted in a two-level nested manner. In the first level, the APC algorithm was applied to all 140 drugs. The pair-wise similarity matrix required by the algorithm as input data was calculated from the Euclidean distance between $\hat{\mathbf{Z}}_a$ and $\hat{\mathbf{Z}}_b$ where a and b refer to two unique drugs. In the second level, the APC algorithm was re-applied to each first-level drug group. Coefficient estimates $\hat{\mathbf{B}}_k^c$ for $k = 1, \ldots, 3$ were used to calculate the pair-wise similarities. Consequently, drugs having similar cell line assignments and cluster-wise coefficient estimates were grouped together. Resulting drug groups are shown in Web Figure 6, where the coloring and thickness of connecting lines indicate degree of closeness. In the third step, each second-level drug group C from step 2 was fitted by a generalized FMMR model, from which we got drug-group-specific cell line assignments, $\hat{\mathbf{Z}}_C$, and cluster-wise coefficient estimates $\hat{\mathbf{B}}_k^C$ for $k = 1, \ldots, 3$.

As an illustration, cell line members and coefficient estimates of clusters 1 and 12 for each second-level drug group are shown in Web Figures 7–8. Cell line members of a cluster are depicted via frequencies of the cancer types cell lines belong to. The frequencies are represented by the size of colored bubbles in top panel of the figures. Estimates of other clusters are not presented in the paper due to limited space. It's very clear that clusters being driven by different cancer types have very different therapeutic genomic profiles and yet no clusters are homogeneous in a particular cancer type. This speaks to the great heterogeneity seen in cancer overall where in fact, cancer is not thought of as a single disease but rather many diseases characterized by different underlying biological changes.

Direct Q2: Can drug resistance patterns be identified along with predictive

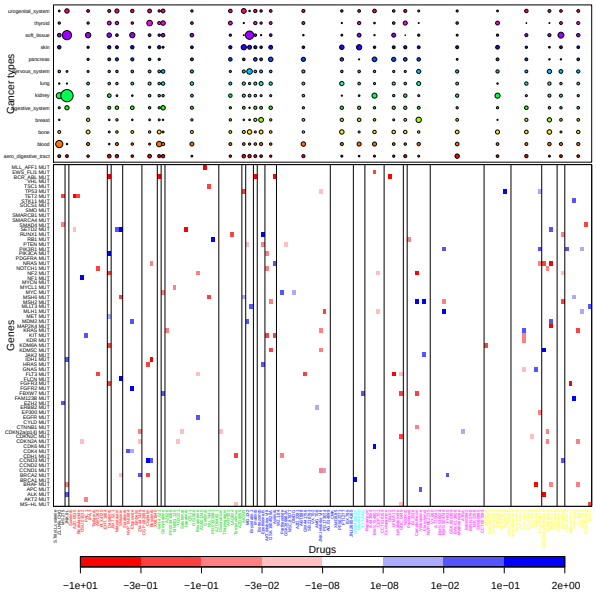

Figure 2: Subset estimation results of cluster 1 by fitting the data of each second-level drug group in a generalized FMMR model. Bottom panel shows a subset of the coefficient estimates corresponding to the mutations of 71 cancer genes, and upper panel shows cell line compositions of cluster 1 by cancer types.

strategies to circumvent resistance using alternative drugs?

We discussed in the Introduction to the paper the success story of Gleevec which was used for the treatment of CML based on the known specificity of targeting the BCR-ABL gene. It is now becoming more and more common in the treatment of cancer to first sequence known cancer genes in tumor genomes and then design therapies accordingly (Bailey et al., 2014). It is also true that tumors can develop resistance to first line therapies by accumulating mutations which confer resistance. We now show how our methodology can be used to identify predictive strategies for circumventing drug resistance based on mutation data.

Figures 2–3 are similar to Web Figures 7–8, except that they only show a subset of $\hat{\mathbf{B}}_1^C$ and $\hat{\mathbf{B}}_{12}^C$ with regard to the mutations of 71 cancer genes, where $\hat{\mathbf{B}}_{12}^C = \hat{\mathbf{B}}_1^C + \hat{\mathbf{B}}_2^C$ are coefficient estimates of cluster 12 for drug group C. Bottom panel of figure 2–3 are in fact regression based drug-genetic association map with red indicating drug-sensitivity biomarkers and blue indicating drug-resistance biomarkers, conditioning on a cluster of cell lines (upper panel) identified by the generalized FMMR model. We propose that above estimation results can be used for drug repurposing (Martins et al., 2015) for resistant tumors. Take soft tissue cancers as an example of this drug repurposing. We extracted those clusters which contain a high percentage of soft tissue cancers (large purple bubbles), along with the subset cluster-wise coefficient estimates. The extracted information is presented in Web Figure 5. Its upper panel shows the IC_{50} values of soft tissue cell lines in extracted clus-

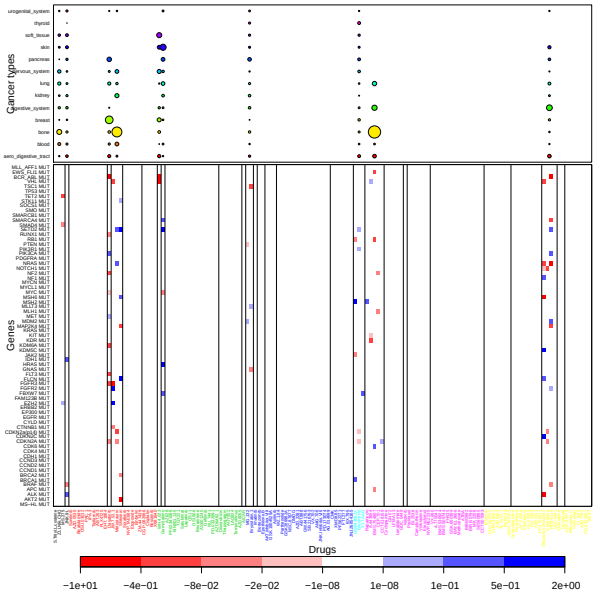

Figure 3: Subset estimation results of cluster 12 by fitting the data of each second-level drug group in a generalized FMMR model.

ters. The drugs in Web Figure 5 were further filtered by only keeping those that have low IC_{50} values without resistance biomarkers and that have high IC_{50} values with resistance biomarkers, leading to Figure 4. In this figure, we conclude that the former set of drugs can be used as alternatives for the later set of drugs which progress resistance to soft tissue cancers. As a positive control of this strategy, drug Gemcitabine was found to be effective in soft tissue cancers resistant to standard chemotherapy (doxyrubicin) in Merimsky et al. (2000).

Direct Q3: Can drug synergies be predicted?

Our desire to allow overlapping cell line clusters becomes apparent in answering this question. Overlapping clusters can be used to guide drug combinations and identify potential drug synergies. From $\hat{\mathbf{Z}}_C$, we can get the cluster-wise IC_{50} values for each drug. Web Figure 9 are boxplots of the cluster-wise IC_{50} values for 15 out of 140 drugs. Due to limited space, the rest are not shown. We then identify the drugs for which the magnitudes of IC_{50}s of overlapping cluster, say 12, are between those of clusters 1 and 2. For instance, drugs MS.275 and GW843682X show such interesting pattern. We focus on the cluster with highest IC_{50}s on average, say cluster 2, and use the coefficient estimates of overlapping cluster to guide the search of another drug to decrease the IC_{50}s of cluster 2 towards overlapping cluster.

Take GW843682X for an example, cluster 2 has the highest IC_{50}s among clusters 1, 2, and 12. We predicted the IC_{50}s in cluster 2 using $\hat{\mathbf{B}}_1^{gw843682x}$, $\hat{\mathbf{B}}_2^{gw843682x}$ and $\hat{\mathbf{B}}_{12}^{gw843682x}$ respectively. Predicted results are shown in top panel (middle) of Figure 5. By doing so,

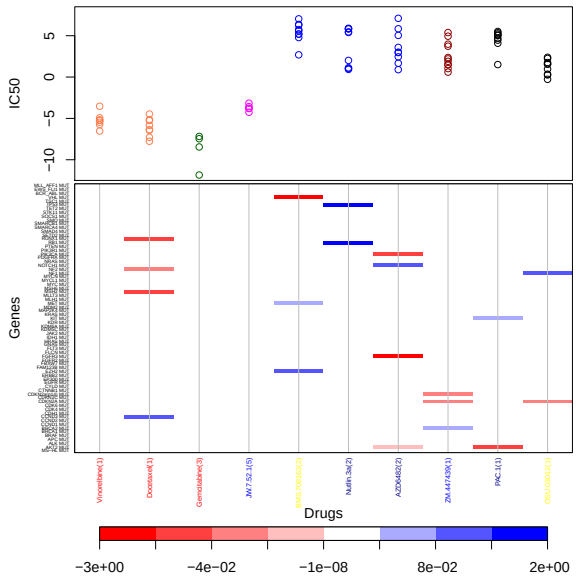

Figure 4: The extracted IC_{50} values of soft tissue cancers (upper panel) and a subset of the coefficient estimates corresponding to the mutations of 71 cancer genes (bottom panel) after filtering those drugs not showing the desired pattern.

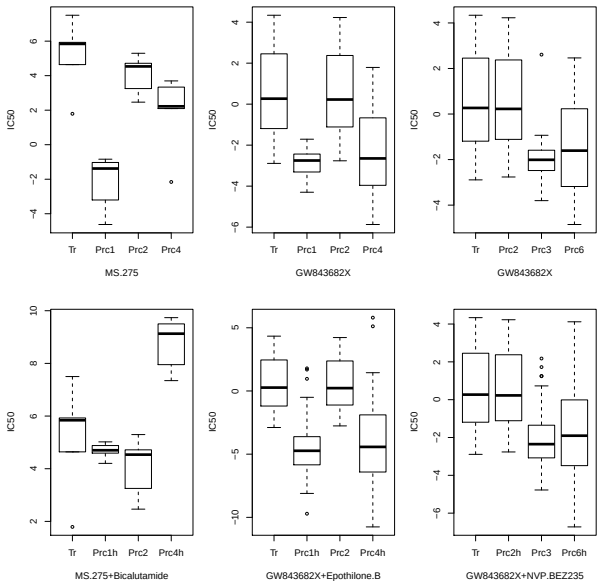

Figure 5: Prediction of drug effects and drug combination effects. In each plot, "Tr" represents the observed IC_{50} values; the rest three boxplots are predicted IC_{50} values using selected coefficient matrices.

we want to see if coefficients $\hat{\mathbf{B}}_1^{gw843682x}$ or $\hat{\mathbf{B}}_{12}^{gw843682x}$ applied to cell lines in cluster 2 can generate low IC_{50}s similar to those in clusters 1 or 12. Since $\hat{\mathbf{B}}_1^{gw843682x}$ corresponds to the coefficient estimates of the target drug GW843682X, we need to find another drug producing similar coefficient estimates to $\hat{\mathbf{B}}_1^{gw843682x}$, this leads to $\hat{\mathbf{B}}_3^{epothilone.b}$. The predicted IC_{50}s of cluster 2 using $\hat{\mathbf{B}}_3^{epothilone.b}$, $\hat{\mathbf{B}}_2^{gw843682x}$ and $\hat{\mathbf{B}}_3^{epothilone.b}+\hat{\mathbf{B}}_2^{gw843682x}$ are presented in bottom panel (middle) of Figure 5. It shows that the combination of GW843682X and Epothilone.B produces much lower IC_{50}s for cluster 2 than GW843682X alone. Similar studies were conducted for MS.275 with respect to its clusters 1, 2, 12 and GW843682X with respect to its clusters 2, 3, 23. Results are shown in left and right columns of Figure 5 respectively. The combination of GW843682X and NVP.BEZ235 decreases the IC_{50}s dramatically, while the combination of MS.275 and Bicalutamide does not show such effect. As a positive control, Wildey et al. (2014) pointed out that a group of drugs (including GW843682X and Epothilone B) may provide a practical starting point to investigate combinatorial drug therapies for synergistic effect in small-cell lung cancer (SCLC). Among the group of drugs, they found synergism between GW843682X and CGP60474 via a preliminary study.

5 Discussion

In this paper, we proposed a new model for identifying therapeutic biomarkers for cancer which can answer specific questions regarding sensitivity, resistance and synergy. We used a penalized likelihood approach for the FMMR model which enforced sparsity in the genomic features. To enable overlapping clustering, the FMMR model was then generalized and a new EM algorithm derived for estimation of model parameters. While the noteworthy plaid model can also be generalized for overlapping clustering multivariate regression data, the generalized FMMR model markedly outperforms this method.

Some improvements can be made for future developments on this work. First, other useful penalty functions for multivariate regression estimations can be adopted, such as the MAP (MAster Predictor) penalty in Peng et al. (2010), the L_2SVS method by Similä and Tikka (2007) and the $L_\infty SVS$ method by Turlach et al. (2005). Second, the multivariate normal distribution assumption on $\mathbf{Y}_i$ can be extended to other more flexible parametric families of multivariate distributions, such as the skew-normal distribution (Azzalini, 2005) and the multivariate skew-t distribution (Chen et al., 2014), for frequent presence of skewness and kurtosis in real data.

As described in the CGP data analyses, we had about 37% of the final analyses observations that had missing data and which we dealt with by using random forest imputation. It would be of interest to explore the impact of this step more thoroughly. An alternative would be to generalize to this problem, the recently developed E-MS algorithm Jiang et al. (2015) for model selection with incomplete data.

Acknowledgements

Hongmei Liu was supported by a doctoral fellowship from the Sylvester Comprehensive Cancer Center at the University of Miami.

Supplementary Materials

Web Appendices, Tables, and Figures referenced in Sections 2.6, 3 and 4 are available with this paper at the Biometrics website on Wiley Online Library.

References

H. Akaike. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 1974.

A. Azzalini. The skew-normal distribution and related multivariate families. *Scandinavian Journal of Statistics*, 32(2):159–188, 2005.

R. Baeza-Yates, B. Ribeiro-Neto, et al. Modern information retrieval. 463, 1999.

A. M. Bailey, Y. Mao, J. Zeng, V. Holla, A. Johnson, L. Brusco, K. Chen, J. Mendelsohn, M. J. Routbort, G. B. Mills, et al. Implementation of biomarker-driven cancer therapy: existing tools and remaining gaps. *Discovery medicine*, 17(92):101, 2014.

A. Banerjee, C. Krumpelman, J. Ghosh, S. Basu, and R. J. Mooney. Model-based overlapping clustering. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 532–537. ACM, 2005.

L. Breiman and J. H. Friedman. Predicting multivariate responses in multiple linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(1): 3–54, 1997.

L. Chen, M. Pourahmadi, and M. Maadooliat. Regularized multivariate regression models with skew-t error distributions. *Journal of Statistical Planning and Inference*, 149:125–139, 2014.

G. Claeskens, N. L. Hjort, et al. Model selection and model averaging. 330, 2008.

A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.

W. S. DeSarbo and W. L. Cron. A maximum likelihood methodology for clusterwise linear regression. *Journal of classification*, 5(2):249–282, 1988.

B. J. Druker, F. Guilhot, S. G. O'Brien, I. Gathmann, H. Kantarjian, N. Gattermann, M. W. Deininger, R. T. Silver, J. M. Goldman, R. M. Stone, et al. Five-year follow-up of patients receiving imatinib for chronic myeloid leukemia. *New England Journal of Medicine*, 355(23):2408–2417, 2006.

J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.

B. J. Frey and D. Dueck. Clustering by passing messages between data points. *science*, 315 (5814):972–976, 2007.

Q. Fu and A. Banerjee. Multiplicative mixture models for overlapping clustering. In *2008 Eighth IEEE International Conference on Data Mining*, pages 791–796. IEEE, 2008.

M. J. Garnett, E. J. Edelman, S. J. Heidorn, C. D. Greenman, A. Dastur, K. W. Lau, P. Greninger, I. R. Thompson, X. Luo, J. Soares, et al. Systematic identification of genomic markers of drug sensitivity in cancer cells. *Nature*, 483(7391):570–575, 2012.

B. Grun and F. Leisch. Flexmix version 2: finite mixtures with concomitant variables and varying and constant parameters. 2008.

H. Ishwaran, U. B. Kogalur, E. H. Blackstone, and M. S. Lauer. Random survival forests. *The annals of applied statistics*, pages 841–860, 2008.

K. Jedidi, V. Ramaswamy, W. S. DeSarbo, and M. Wedel. On estimating finite mixtures of multivariate regression and simultaneous equation models. *Structural Equation Modeling: A Multidisciplinary Journal*, 3(3):266–289, 1996.

J. Jiang, T. Nguyen, and J. S. Rao. The e-ms algorithm: model selection with incomplete data. *Journal of the American Statistical Association*, 110(511):1136–1147, 2015.

P. Jones and G. McLachlan. Fitting finite mixture models in a regression context. *Australian Journal of Statistics*, 34(2):233–240, 1992.

C. Keribin. Consistent estimation of the order of mixture models. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 49–66, 2000.

A. Khalili and J. Chen. Variable selection in finite mixture of regression models. *Journal of the american Statistical association*, 2012.

A. Khalili and S. Lin. Regularization in finite mixture of regression models with diverging number of parameters. *Biometrics*, 69(2):436–446, 2013.

L. Lazzeroni and A. Owen. Plaid models for gene expression data. *Statistica sinica*, pages 61–86, 2002.

M. M. Martins, A. Y. Zhou, A. Corella, D. Horiuchi, C. Yau, T. Rakshandehroo, J. D. Gordan, R. S. Levin, J. Johnson, J. Jascur, et al. Linking tumor mutations to drug responses via a quantitative chemical–genetic interaction map. *Cancer discovery*, 5(2): 154–167, 2015.

O. Merimsky, I. Meller, G. Flusser, Y. Kollender, J. Issakov, M. Weil-Ben-Arush, E. Fenig, G. Neuman, D. Sapir, S. Ariad, et al. Gemcitabine in soft tissue or bone sarcoma resistant to standard chemotherapy: a phase ii study. *Cancer chemotherapy and pharmacology*, 45 (2):177–181, 2000.

R. Needham. Computer methods for classification and grouping. *The Use of Computers in Anthropology, I. Hymes, ed*, pages 345–356, 1965.

J. Peng, J. Zhu, A. Bergamaschi, W. Han, D.-Y. Noh, J. R. Pollack, and P. Wang. Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer. *The annals of applied statistics*, 4(1):53, 2010.

R. E. Quandt. A new approach to estimating switching regressions. *Journal of the American statistical association*, 67(338):306–310, 1972.

A. J. Rothman, E. Levina, and J. Zhu. Sparse multivariate regression with covariance estimation. *Journal of Computational and Graphical Statistics*, 19(4):947–962, 2010.

G. Schwarz et al. Estimating the dimension of a model. *The annals of statistics*, 6(2): 461–464, 1978.

T. Similä and J. Tikka. Input selection and shrinkage in multiresponse linear regression. *Computational Statistics & Data Analysis*, 52(1):406–422, 2007.

M. Stone. An asymptotic equivalence of choice of model by cross-validation and akaike's criterion. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 44–47, 1977.

R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

B. A. Turlach, W. N. Venables, and S. J. Wright. Simultaneous variable selection. *Technometrics*, 47(3):349–363, 2005.

H. Turner, T. Bailey, and W. Krzanowski. Improved biclustering of microarray data demonstrated through systematic performance tests. *Computational statistics & data analysis*, 48(2):235–254, 2005.

G. Wildey, Y. Chen, I. Lent, L. Stetson, J. Pink, J. S. Barnholtz-Sloan, and A. Dowlati. Pharmacogenomic approach to identify drug sensitivity in small-cell lung cancer. *PloS one*, 9(9):e106784, 2014.

J. Zhang. A bayesian model for biclustering with applications. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 59(4):635–656, 2010.

H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.

Web-based Supplementary Materials for "Precision Therapeutic Biomarker Identification with Application to the Cancer Genome Project"

A Simulation Studies

A.1 Design of the Simulations

Two scenarios were designed to investigate the non-overlapping clustering and overlapping clustering of the new method respectively. In each scenario, data were generated from model (2.1) with $K = 3$, $p_n = 15$, $q = 3$ and $n = 150, 450$. Predictors $\mathbf{x}_i$, $i = 1, \ldots, n$, were drawn independently from $N_{p_n}(\mathbf{0}, \mathbf{\Sigma}_1)$ with $\mathbf{\Sigma}_1(i, j) = 0.5^{|i-j|}$. Random errors ε_i, $i = 1, \ldots, n$, were drawn independently from $N_q(\mathbf{0}, \mathbf{\Sigma}_2)$ with $\mathbf{\Sigma}_2(i, j) = 0.75^{|i-j|}$.

The sparse coefficient matrices $\mathbf{B}_k, k = 1, 2, 3$, were generated as

$$\mathbf{B}_k = \mathbf{W} \otimes \mathbf{S} \otimes \mathbf{T}, \tag{18}$$

where $\otimes$ indicates the element-wise product. Each entry of $\mathbf{W}$ was drawn independently from $N(0, 1)$, each entry of $\mathbf{S}$ was drawn independently from the Bernoulli distribution $B(1, p_1)$, and each row of $\mathbf{T}$ (either all 1 or all 0) was determined by an independent draw from $B(1, p_2)$. As a result, we expect $p_1 p_2 p_n$ relevant predictors for each response variable, and $(1 - p_2)p_n$ predictors are expected to be irrelevant to all q response variables. We let $p_1 = 0.5$ and $p_2 = 0.9$. Each simulation was repeated for 50 times. A new sequence of $\mathbf{B}_k, k = 1, 2, 3$ and a new set of data $(\mathbf{X}, \mathbf{Y})$ were generated for each repetition. We assumed that K is known, so that performance evaluation of the new method is not interfered by selection of K. Its selection is separately examined in Section A.4.

Scenario 1: The 3 clusters are non-overlapping, each cluster contains $n/3$ observations.

Scenario 2: 70% of the observations involve in single cluster, 22% of the observations involve in two clusters and 8% of the observations involve in three clusters.

A.2 Quality Measures of Cluster Recovery

Quality measures in Baeza-Yates et al. (1999) have been used to evaluate the clustering performance of the proposed method. Denote N_S the number of elements in S. Below quality measures were proposed to compare a retrieved cluster $\hat{A}$ with a target cluster A,

$$\text{"specificity"} = \frac{N_{A\cap\hat{A}}}{N_A}, \quad \text{"sensitivity"} = \frac{N_{A\cap\hat{A}}}{N_{\hat{A}}},$$
$$\text{"}F_1 \text{ measure"} = \frac{2N_{A\cap\hat{A}}}{N_A + N_{\hat{A}}}.$$

The F_1 measure, taken as the harmonic mean of specificity and sensitivity, gives an overall measure of the clustering.

Moreover, we employed the one-to-one correspondence match approach (Turner et al., 2005) to evaluate a sequence of retrieved clusters. It includes about three steps. First make the number of retrieved clusters to be the same as the number of target clusters by adding

null clusters to retrieved clusters or dropping addition poorly retrieved clusters. Retrieved clusters are then matched to target clusters in pair. Finally, calculate the mean quality measures for the sequence of paired clusters.

A.3 Results of the Simulations

A sequence of methods were implemented for comprehensive comparison with our new method, including the generalized plaid model (*plaid*) and its revised counterparts (*aplaid*); the R package *flexmix* (*EM*) where response variables are treated as independent and overlap issues are not considered; the generalized FMMR model with response variables treated as independent (*gEMseplasso0*); the generalized FMMR model with separate LASSO estimation of $\mathbf{B}_k$ (*gEMseplasso*); and finally the original generalized FMMR model with MRCE estimation of $\mathbf{B}_k$ (*gEMmrce*).

Their performances were evaluated via the three quality measures introduced in Section A.2 and the sum of squared errors (SSE) of coefficients estimates. Results are summarized in Web Figure 6–7. In scenario 1 when there is no overlap, *EM* and *gEMseplasso0* do equally well showing that the generalized FMMR model retains the ability to recover non-overlapping clusters. By taking into account of the covariance structure among response variables, *gEMseplasso* and *gEMmrce* outperform *EM* and *gEMseplasso0*, but there is little difference between *gEMseplasso* and *gEMmrce*. In scenario 2 when there are 30% overlaps, specificity of *EM* drops down dramatically, resulting in a poor F_1 measure. The *gEMseplasso0* does much better than *EM* by estimating the overlaps, while it is defeated by *gEMseplasso* and *gEMmrce*. In this scenario, *gEMmrce* outperforms *gEMseplasso* when $n = 150$. However, they converge very fast with n, both achieving a median 95% clustering accuracy when $n = 450$. The *plaid* and *aplaid* have a poor performance in all scenarios.

A.4 Additional Simulations

We conducted another simulation for scenario 2 to evaluate the accuracy of BIC in selection of K. In the simulation, K is unknown in candidate procedures *EM* and *gEMseplasso*, and is chosen by minimizing the BIC. Results are summarized in Web Figure 8 and Web Table 1. Given 30% overlaps in simulated data, BIC has a poor performance in identifying the true K with the *EM* procedure. Whereas by fitting the overlapping clusters with *gEMseplasso*, BIC obtains a much higher accuracy in selection of K (see Web Table 1). Moreover, although K is misspecified for 12% of the times when $n = 450$, *gEMseplasso* still has achieved a median 93% clustering accuracy (see Web Figure 8).

An additional simulation was conducted for scenario 2 to investigate the advantage of modeling with multivariate response variables in the new method. We compared the performance of modeling with all 3 response variables versus modeling with each of the 3 response variables in the generalized FMMR model. Results are summarized in Web Figure 9. Using multivariate response variables leads to a much higher clustering accuracy.

Algorithm 1 Plaid

For $k = 1 : K$ do

At this point, we have already found $k-1$ layers and are searching for the kth layer. Calculate the residuals from previous $k-1$ layers by

$$\mathbf{z}_i = \mathbf{y}_i - \sum_{l=1}^{k-1} \mathbf{x}_i^T \mathbf{B}_l P_{il}, \quad i = 1, \ldots, n.$$

Initialize $\mathbf{P}_k^0$.

For $s = 1 : S$ do

Given $\mathbf{P}_k^{s-1}$, compute

$$\mathbf{B}_k^s = \arg\min_{\mathbf{B}_k} \frac{1}{2} \sum_{i=1}^{n} \|\mathbf{z}_i - \mathbf{x}_i^T \mathbf{B}_k P_{ik}^{s-1}\|_2^2 + \sum_{j=1}^{q} \lambda_{jk} \|\mathbf{b}_{jk}\|_1,$$

where $\mathbf{B}_k = (\mathbf{b}_{1k}, \ldots, \mathbf{b}_{qk})$.

Given $\mathbf{B}_k^s$, compute

$$P_{ik}^s = \begin{cases} 0.5 + \min\{0.5, \frac{S}{2(S-T)}\}, & \|\mathbf{z}_i - \mathbf{x}_i^T \mathbf{B}_k^s\|_2^2 \leq \|\mathbf{z}_i\|_2^2, \\ 0.5 - \min\{0.5, \frac{S}{2(S-T)}\}, & \text{otherwise,} \end{cases}$$

where T was taken as $0.2S$ for $i = 1, \ldots, n$.

End for

Given $\mathbf{P}_k^S$, update $\mathbf{B}_k^{S+1}$.

Given $\mathbf{B}_k^{S+1}$, prune layer memberships by

$$P_{ik}^{S+1} = \begin{cases} 1, & \|\mathbf{z}_i - \mathbf{x}_i^T \mathbf{B}_k\|_2^2 \leq \tau \|\mathbf{z}_i\|_2^2, \\ 0, & \text{otherwise,} \end{cases}$$

where $\tau \in (0, 1)$ for $i = 1, \ldots, n$.

Given $\mathbf{P}_k^{S+1}$, update $\mathbf{B}_k^{S+2}$.

Given $\mathbf{P}_1^{S+1}$, ..., $\mathbf{P}_k^{S+1}$, back fit the layers R times. We set $R = 2$.

End for

Algorithm 2 All-plaid

Initialize $\mathbf{B}_1^0, \ldots, \mathbf{B}_K^0$ and $\mathbf{P}_1^0, \ldots, \mathbf{P}_K^0$.
For $s = 1 : S$ do
For $k = 1 : K$ do
Given $(\mathbf{B}_1^s, \ldots, \mathbf{B}_{k-1}^s, \mathbf{B}_{k+1}^{s-1}, \ldots, \mathbf{B}_K^{k-1})$ and $(\mathbf{P}_1^s, \ldots, \mathbf{P}_{k-1}^s, \mathbf{P}_k^{s-1}, \ldots, \mathbf{P}_K^{s-1})$, compute

$$\mathbf{z}_i = \mathbf{y}_i - \sum_{l=1}^{k-1} \mathbf{x}_i^T \mathbf{B}_l^s P_{il}^s - \sum_{l=k+1}^{K} \mathbf{x}_i^T \mathbf{B}_l^{s-1} P_{il}^{s-1}, \quad i = 1, \ldots, n,$$

$$\mathbf{B}_k^s = \arg\min_{\mathbf{B}_k} \frac{1}{2} \sum_{i=1}^{n} \|\mathbf{z}_i - \mathbf{x}_i^T \mathbf{B}_k P_{ik}^{s-1}\|_2^2 + \sum_{j=1}^{q} \lambda_{jk} \|\mathbf{b}_{jk}\|_1,$$

where $\mathbf{B}_k = (\mathbf{b}_{1k}, \ldots, \mathbf{b}_{qk})$.
Given $\mathbf{B}_k^s$, compute

$$P_{ik}^s = \begin{cases} 0.5 + \min\{0.5, \frac{S}{2(S-T)}\}, & \|\mathbf{z}_i - \mathbf{x}_i^T \mathbf{B}_k\|_2^2 \leq \tau \|\mathbf{z}_i\|_2^2, \\ 0.5 - \min\{0.5, \frac{S}{2(S-T)}\}, & \text{otherwise}, \end{cases}$$

where T was taken as $0.2S$ and $\tau \in (0, 1)$ for $i = 1, \ldots, n$.
End for
End for

Web Table 1: Accuracy of BIC in selection of K.

sample size	*EM*	*gEMseplasso*
n=150	0.32	0.56
n=450	0.22	0.88

B Two algorithms of the Generalized Plaid Model

Web Tables

Web Figures

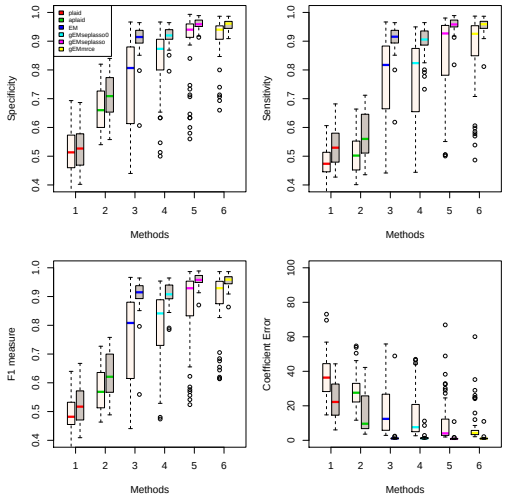

Web Figure 6: Simulation results of scenario 1. The two box plots under each method correspond to $n = 150, 450$ respectively.

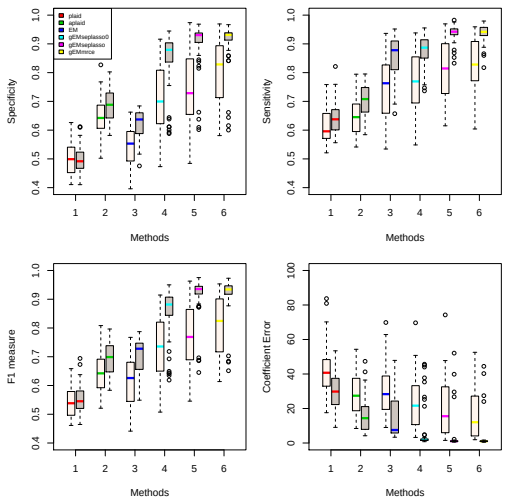

Web Figure 7: Simulation results of scenario 2.

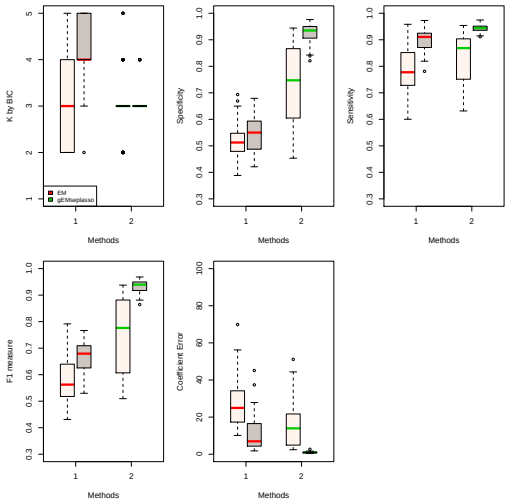

Web Figure 8: Simulation results of BIC in selection of K.

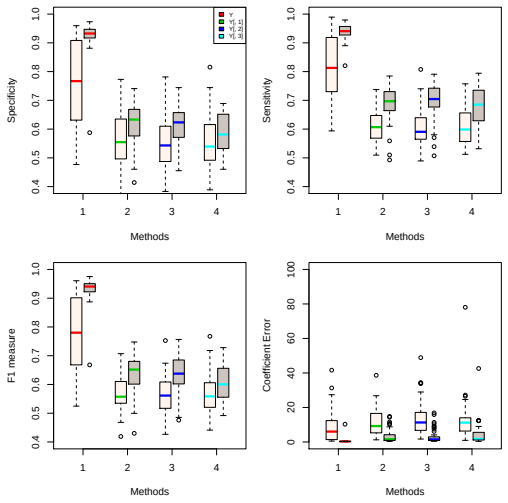

Web Figure 9: Simulation results of using multivariate responses versus univariate responses.

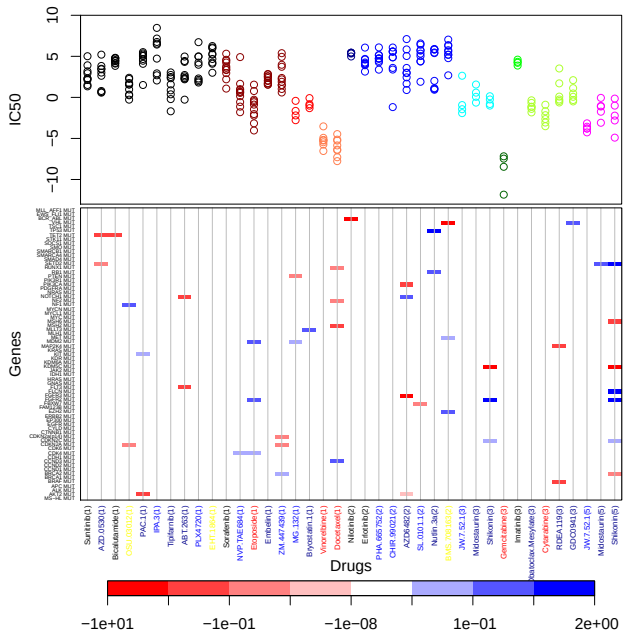

Web Figure 10: The extracted IC_{50} values of soft tissue cancers (upper panel) and a subset of the coefficient estimates corresponding to the mutations of 71 cancer genes (bottom panel). Drug names in red indicate type "chemo", black indicate type "clinical", blue indicate type "experimental", darkblue indicate type "in clinical development" and yellow indicate drug type not available.

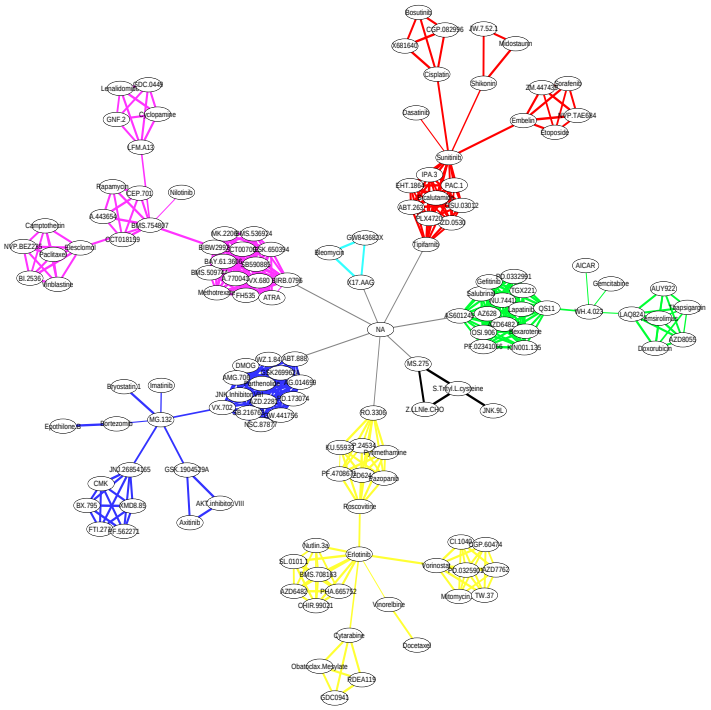

Web Figure 11: Groups of drugs generated from the APC algorithm. Each node implies a type of drug with the drug name on it. Drugs connected with the same color of strokes belong to a first-level drug group, and drugs connected within a closure belong to a second-level drug group.

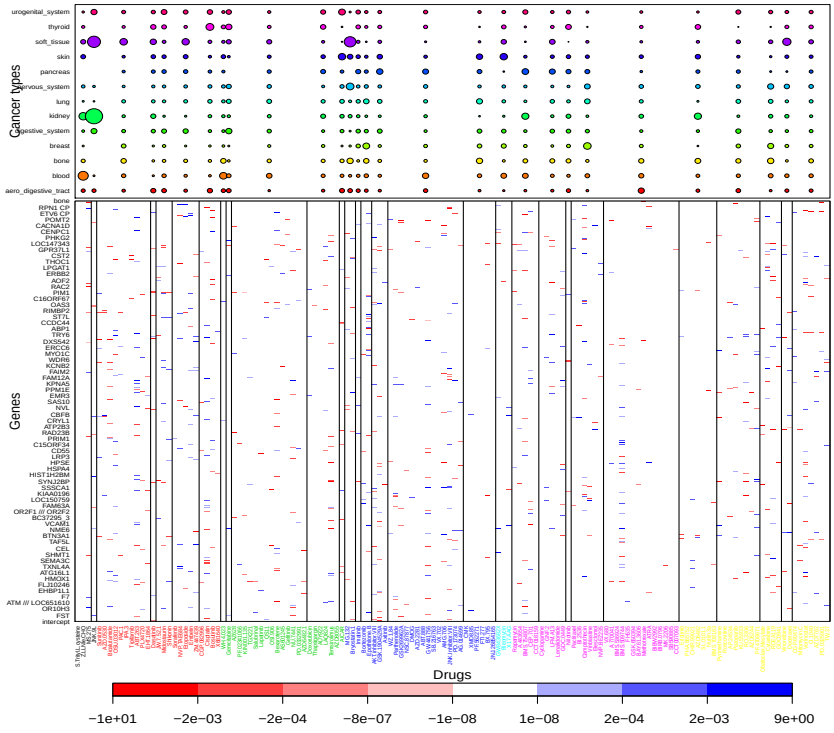

Web Figure 12: Estimation results of cluster 1 by fitting the data of each second-level drug group in a generalized FMMR model. Bottom panel plots the coefficient estimates, and upper panel describes the cell line compositions of cluster 1 by cancer types.

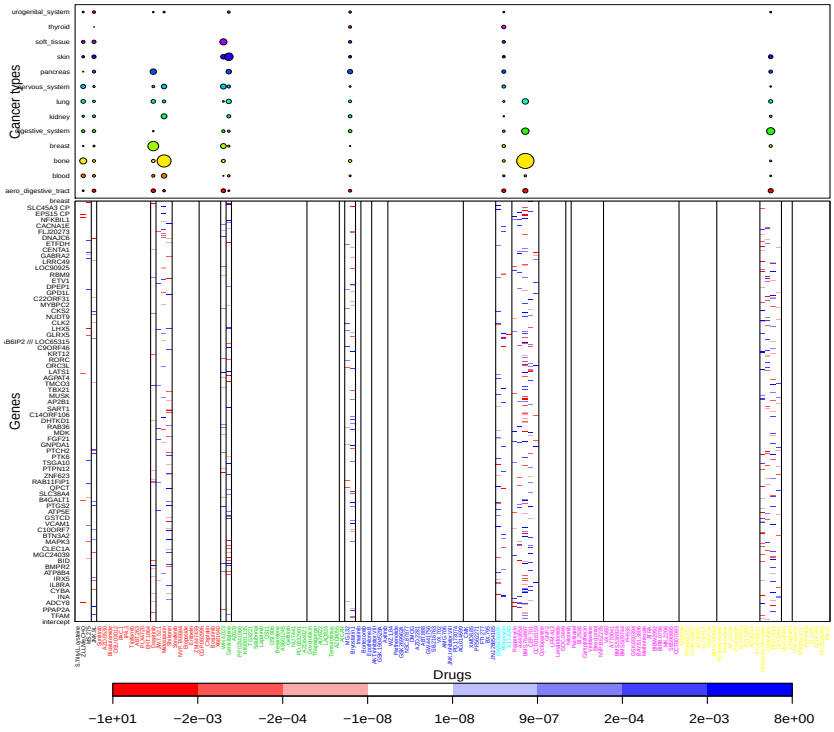

Web Figure 13: Estimation results of cluster 12 by fitting the data of each second-level drug group in a generalized FMMR model.

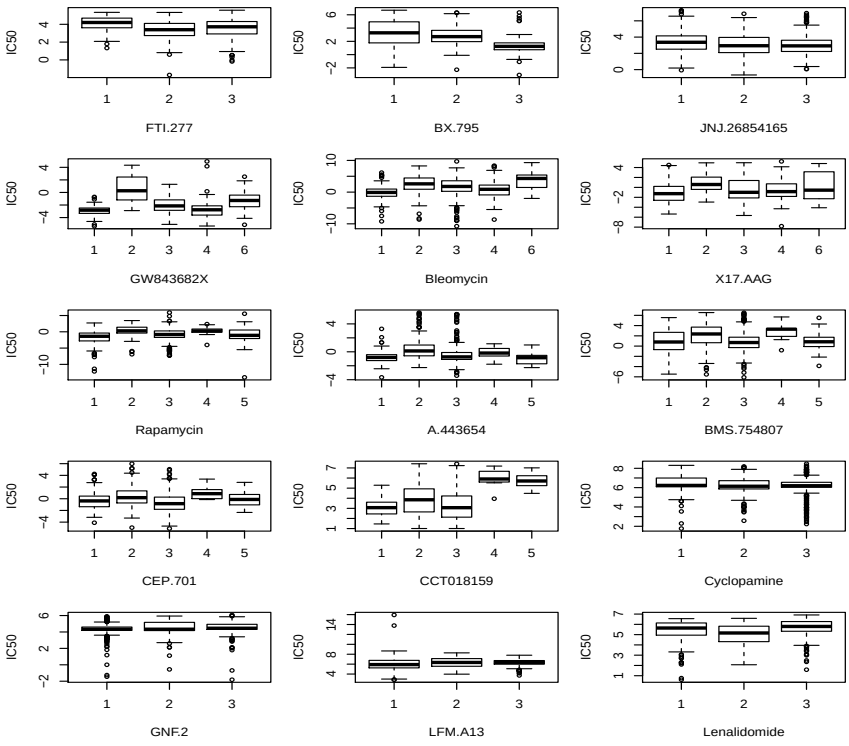

Web Figure 14: Cluster-wise IC_{50} values of each drug.